embodiment the transaction sequence numbers are consecutive and monotonic for all completed transactions.

Although the invention is illustrated with respect to synchronization of two computers **36**, **38**, those of skill in the art will appreciate that a more than two computers may also be synchronized. In such cases, additional transferring steps transfer the transaction updates to other computers **28** in the network **10**, and additional applying steps apply the transaction updates to replicas **56** on the other computers **28**.

In summary the present invention provides a system and method for properly synchronizing transactions when a disconnectable computer **28** is reconnected to the network **10**. The invention is not limited to file system operations but can instead be extended to support a variety of database objects by using the schema **84**, object distributor **82**, object processor **86**, and other modules. Clash handling means may be used to identify potentially conflicting database changes and allow their resolution by either automatic or manual means. Clash handling and retries also make locks optional.

The synchronization time period does not depend directly on the total number of files, directories, or other objects in the replica **56**. Rather, the required time depends on the number and size of the objects that require updating. This facilitates synchronization over slow links, such as mobile computer modems and WAN server connections. Unlike conventional systems such as state-based systems, the present invention is therefore readily scaled upward to handle larger networks. Moreover, the dynamic intervention and processing of operations by the replica managers **46** allows systems of the present invention to support continuous synchronization across slow links. The replica managers **46** also allow the use of consistent file locations regardless of whether a mobile computer is connected to the network **10**.

Although particular methods embodying the present invention are expressly illustrated and described herein, it will be appreciated that apparatus and article embodiments may be formed according to methods of the present invention. Unless otherwise expressly indicated, the description herein of methods of the present invention therefore extends to corresponding apparatus and articles, and the description of apparatus and articles of the present invention extends likewise to corresponding methods.

The invention may be embodied in other specific forms without departing from its essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. Any explanations provided herein of the scientific principles employed in the present invention are illustrative only. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed and desired to be secured by patent is:

1. A method for synchronizing transactions in a system, the system including at least two computers capable of being connected by a network link, each of the computers including a storage device containing a distributed hierarchical database replica, a device controller in signal communication with the storage device, a replica manager in signal communication with the device controller and the network link, and a database manager in signal communication with the replica manager, the method comprising the computer-implemented steps of:

routing database transactions with the database managers through the replica managers to the device controllers;

connecting at least two computers with the network link; and

using each of at least two replica managers to route a transaction to another replica manager after the two computers are connected by the network link.

2. The method of claim **1**, further comprising the computer-implemented step of recording a transaction on a non-volatile storage medium.

3. The method of claim **2**, wherein said recording step comprises the step of entering a representation of the transaction in a transaction log.

4. The method of claim **1** for synchronizing transactions in the system, the system containing a network of connectable computers, the transactions targeting entries in the distributed hierarchical database replicas, the replicas residing on separate computers in the network, said method comprising the steps of:

obtaining a network connection between a first computer and a second computer as part of said step of connecting at least two computers with the network link;

identifying a first transaction that targets a distributed hierarchical database entry in a first replica on the first computer, as part of said step of routing database transactions; and

as part of said step of using each of at least two replica managers, locating a corresponding second replica that resides on the second computer, transferring an update based on the first transaction over the network connection from the first computer to the second computer, and applying the first transaction update to the second replica.

5. The method of claim **4**, wherein the replicas contain file descriptors and directory descriptors for a file system.

6. The method of claim **4**, wherein the replicas contain directory services entries.

7. The method of claim **4**, wherein said step of transferring an update comprises transferring a transaction sequence number corresponding to the first transaction and a location identifier corresponding to the first computer.

8. The method of claim **7**, wherein the first transaction is one of a plurality of transactions completed at the first computer and each completed transaction has a corresponding transaction sequence number.

9. The method of claim **8**, wherein the transaction sequence numbers are generated in a predetermined order.

10. The method of claim **9**, wherein the transaction sequence numbers are consecutive and monotonic for all completed transactions.

11. The method of claim **8**, further comprising the computer-implemented step of detecting a missed update by detecting a gap in a plurality of transferred transaction sequence numbers.

12. The method of claim **7**, wherein the first transaction is one of a plurality of transactions completed at the first computer, each completed transaction has a corresponding transaction sequence number, and the transaction sequence numbers are consecutive and monotonic for all completed transactions.

13. The method of claim **12**, further comprising the computer-implemented step of detecting a missed update by detecting a gap in a plurality of transferred transaction sequence numbers.

14. The method of claim **4**, wherein said transferring step further comprises transferring the first transaction update to at least one computer other than the first and second computers, and said applying step further comprises applying the first transaction update to at least one replica other than the first and second replicas.

15. The method of claim **4**, wherein said applying step comprises setting a database object lock that serializes updates to the first replica.